ROML: A Robust Feature Correspondence Approach for Matching Objects in A Set of Images

Kui Jia · Tsung-Han Chan · Zinan Zeng · Gang Wang · Tianzhu Zhang · Yi Ma

arXiv:1403.7877v1 [cs.CV] 31 Mar 2014

Abstract Feature-based object matching is a fundamental problem for many applications in computer vision, such as object recognition, 3D reconstruction, tracking, and motion segmentation. In this work, we consider simultaneously matching object instances in a set of images, where both inlier and outlier features are extracted. The task is to identify the inlier features and establish their *consistent* correspondences across the image set. This is a challenging combinatorial problem, and the problem complexity grows exponentially with the image number. To this end, we propose a novel framework, termed ROML, to address this problem. ROML optimizes simultaneously a partial permutation matrix (PPM) for each image, and feature correspondences are established by the obtained PPMs. Our contributions are summarized as follows: (1) We formulate the problem as rank and sparsity minimization for PPM optimization, and treat simultaneous optimization of multiple PPMs as a regularized consensus problem in the context of distributed optimization. (2) We use the ADMM method to solve the thus formulated ROML problem, in which a subproblem associated with a single PPM optimization appears to be a difficult integer quadratic program (IQP). We prove that under wildly applicable conditions, this IQP is equivalent to a linear sum assignment problem (LSAP), which can be efficiently solved to an exact solution. (3) Our framework is independent of the domain of applications and type of features. (4) Extensive experiments on object matching, tracking, and localization show the superiority of our method over existing ones.

Keywords Object matching · Feature correspondence · Low-rank · Sparsity.

Kui Jia · Tsung-Han Chan · Zinan Zeng · Tianzhu Zhang
Advanced Digital Sciences Center, 1 Fusionopolis Way, Singapore
E-mail: kuijia@gmail.com; chantsunghan@gmail.com; edwin.zeng@adsc.com.sg; tz.zhang@adsc.com.sg

Gang Wang
School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore
E-mail: wanggang@ntu.edu.sg

Yi Ma
Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA
E-mail: yima@illinois.edu
School of Information Science and Technology, ShanghaiTech University, China
E-mail: mayi@shanghaitech.edu.cn

1 Introduction

Object matching is a fundamental problem in computer vision. Given a pair or a set of images that contain common object instances, or an object captured under varying poses, it involves establishing correspondences between the parts or features of the objects contained in the images. Accurate, robust, and consistent matching across images is a key ingredient in a wide range of applications such as object recognition, shape matching, 3D reconstruction, tracking, and motion segmentation.

For a pair of feature sets extracted from two images, finding inliers from them and establishing correspondences are in general a combinatorial search problem. Objects may appear in images with cluttered background, and some parts of the objects may also be occluded. The search space can further explode when a *globally consistent matching* across a set of images is desired. For object instances with large intra-category variations or those captured under varying poses (e.g., non-rigid objects with articulated pose changes), the matching tasks become even more difficult. All these factors make object matching a very challenging task.

In literature, a variety of strategies have been proposed for object matching. In particular, early shape matching works use point sets to represent object patterns [53,55]. To match between a pair of point sets, they build point descriptions by modeling spatial relations of points within each point set as higher level geometric structures, e.g., lines, curves, and surfaces, or more advanced features, e.g., shape context [5]. In [10, 17,5], alternating estimation of point correspondence and geometric transformation is also used for non-rigid shape matching. In general, point set based shape matching is less robust to measurement noise and outliers, with classical techniques such as RANSAC [30] available to improve its robustness. The development of local invariant features [42,45] for discriminative description of visual appearance has brought significant progress in object matching and recognition [54]. For example in [52,23], instances of a common object category from an image collection can be located and matched by exploiting the discriminative power of local feature descriptors. The popular Bag-of-Words model for object recognition is also built on matching (clustering) similar local region descriptors. However, local descriptors alone can be ambiguous for matching when there exist repetitive textures or less discriminative local appearance in images. In between of these two extremes, recent graph matching methods [37,60,65] consider both feature similarity and geometric compatibility between two sets of features, where the nodes of graphs correspond to local features and edges encode spatial relations between them. Mathematically, graph matching is formulated as a quadratic assignment problem (QAP), which is known to be NP-hard. Although intensive efforts of these methods have been focusing on devising more accurate and efficient algorithms to solve this problem, in general, they can only obtain approximate solutions for QAP, and thus suboptimal correspondences for robust object matching.

Most of these methods focus on establishing correspondences between a pair of images. However, in practice, it is very common that when such a pair of images are available, a set of images are also available that we know a common object is present in them, such as a video sequence with a moving object, or a set of images collected from the Internet that contain instances of a generic object category. In these situations, it is desired that a globally consistent matching can be established. This is a very challenging combinatorial problem. As the number of images increases, the problem complexity explodes exponentially. A straightforward approach is to locally build correspondences between pairs of images. Obviously, pair-wise matching can only get suboptimal solutions, since matching found between pairs of images may not be globally consistent across the whole set. Compared to global matching, pair-wise matching is also less robust to outliers and occlusion of inlier features, as it cannot leverage additional constraints from other images that also contain the same object pattern of interest. In this work, we are thus interested in the following object matching problem.

Problem 1: Given a set of images with both inlier and outlier features extracted from each image, simultaneously identify a given number of inlier features from each image and establish their consistent correspondences across the image set.

In Problem 1, we consider the common scenario in object matching that there is exactly one object instance in each image. The inlier features describe local salient appearance of the object, and the rest of the features are outliers. Depending on applications, the types of features can be chosen as image coordinates, local region descriptors [42,22,7], or combination of them [59]. Under such a setting, each image is naturally represented as a set of feature vectors. If we can re-order the correctly identified inlier features from each image, and concatenate them as a (long) vector, ideally the matrix formed by horizontally arraying these vectors should be low-rank [1]. The underlying rationale is that inlier features, the corresponding ones of which are correlated to each other, repetitively appear in the image set, while outliers just accidently do so. In situations when some inliers are missing, or when there are variations in inlier features of different images (e.g., due to illumination or pose changes), we can decompose this matrix as a combination of low-rank and sparse terms, where the sparse term models those variations or missing inliers.

Motivated by these observations, we propose in this paper a novel and principled framework, termed ROML, for robustly matching objects in a set of images. ROML is formulated to optimize a partial permutation matrix (PPM) for each image by rank and sparsity minimization. PPM is composed of binary entries (cf. (1)). It simultaneously selects inlier features from each image and re-orders them, so that after selection and reordering, the matrix formed by these feature vectors can be decomposed into a low-rank matrix and a sparse matrix. The so formulated ROML problem belongs to a more general class of multi-index assignment problem (MiAP), which is proven to be NP-hard [13]. Exact solution methods are prohibitively slow for practical use. In this work, we treat ROML, more specifically simultaneous optimization of multiple PPMs, as a regularized consensus problem in the context of distributed optimization [9]. We use the Alternating Di-

[1] The form of low-rank matrix using image coordinates as features is different, which we will present in Section 4.1.

rection Method of Multipliers (ADMM) [11] to solve the ROML problem, in which a subproblem associated with a single PPM optimization appears to be a difficult integer quadratic program (IQP). We prove that under widely applicable conditions, this IQP is equivalent to a linear sum assignment problem (LSAP) [13], which can be efficiently solved to an exact solution using the Hungarian algorithm [34]. Our framework is independent of the domain of applications and type of features. Extensive experiments on matching instances of a common object category, non-rigid object tracking, and common object localization show the superiority of our method over existing ones.

A preliminary work of this paper has appeared in [63]. We have made the following new contributions in the present paper.

- Although [63] proposes to optimize a set of PPMs via rank and sparsity minimization for robust feature matching, however, its solution of each PPM optimization is obtained by sequentially solving two costly subproblems: a quadratic program over the continuous-domain relaxation of PPM, followed by a binary integer programming that projects the relaxed PPM into its feasible set. In fact, the second subproblem is irrelevant to the original objective function, and consequently, the thus obtained PPM is only suboptimal. In contrast, we propose in the present paper a new method to solve the PPM optimization and prove that under wildly applicable conditions, the PPM optimization step is equivalent to an LSAP, which can be efficiently solved to an exact solution using the Hungarian algorithm. Extensive experiments in Section 5 show the great advantage of the proposed ROML over the method in [63] in terms of both matching accuracy and efficiency.
- We present mathematical analysis in this paper to show that the proposed ROML formulation belongs to the NP-hard MiAP. We also discuss the suitability of ADMM for approximately solving ROML from the perspective of distributed optimization. These analysis and discussion put ROML in a broader context, which are overlooked in [63].
- In terms of feature types used, [63] considers local region descriptors alone, neglecting the exploitation of spatial relations between interest points in each image, which could be very important for reliable matching when region descriptors alone become less distinguishable between each other. Consequently, only toy feature matching experiments on face images are presented in [63]. Instead, we show in this paper that ROML works with a variety of feature types including image coordinates, region descriptors, or combination of them (cf. Section 4).
- We apply ROML to diverse applications such as rigid object matching, matching object instances of a common category, non-rigid object tracking, and localizing common objects in a set of images. Compared to existing methods, we achieve the best performance in most of these applications.

2 Related Works

There is an intensive literature on object/shape matching between a pair of images [20]. Representative works include shape context [5], graph [37,60] and hyper-graph [35,62,24] matching. In this section, we briefly review several existing methods that use multiple images/point sets for object matching, and also the more general MiAP.

Maciel and Costeira [44] first proposed to use PPM to model both feature matching and outlier rejection in a set of images. They formulated optimization of PPMs as an integer constrained minimization problem. To solve this combinatorial problem, they relaxed both the objective function and integer constraints, resulting in an equivalent concave minimization problem. However, the complexity of concave minimization is still non-polynomial. Moreover, matching criteria used in the cost function of [44] were locally based on pair-wise similarity of features in different images. Instead, our method is based on low-rank and sparse minimization (via convex surrogate functions), whose problem size is polynomial w.r.t. the numbers of features and images, and whose cost function is also globally defined over features in all the images.

Rank constraints have been used in [47,48] for point matching across video frames. They constructed a measurement matrix containing image coordinates of points extracted from a moving rigid object. Motivated by factorization model in shape-from-motion [58], they assumed this measurement matrix was low-rank, and used rank constraints to optimize PPMs for establishing point correspondences across frames. The method in [47] is limited in several aspects: (1) an initial template of point set without outliers is assumed given; (2) every inlier point is required to be visible in all frames; (3) matching across frames is a bootstrapping process - points in a subsequent frame are to be aligned to those of previously matched frames, thus matching errors will inevitably propagate and accumulate; (4) an initial rough estimate of point correspondences for a new frame is assumed given in their algorithm, which may be only valid for slow motion objects. The aspects (1) and (3) have to some extent been alleviated in [48],

but [48] cannot cope with the other limiting aspects. As a globally consistent and robust matching framework, our method has no such limitations. More importantly, we note that the mechanism of rank constraints used in [47,48] is different from that of our method. Methods [47,48] can only apply to matching of rigid objects using image coordinates as features, while our method considers low-rank assumption on a type of generally defined features, which take image coordinates and region descriptors as instances. Consequently, our method is able to apply in more general scenarios, such as matching of objects with non-rigid deformation.

Recently, a low-dimensional embedding method was proposed in [59] for feature matching. Given feature points extracted from each of a set of images, it can learn an embedded feature space, which encodes information of both region descriptors and the geometric structure of points in each image. [59] used k-means clustering in the embedded space for feature matching. As we will show in Section 5, k-means based on Euclidean distances of embedded features is not a good way to establish correspondences. There is no explicit outlier rejection mechanism in [59] either. Compared to [59], our method uses the low-rank and sparse constraints to optimize PPMs, which integrates correspondence and outlier rejection in a single step.

As mentioned in Section 1, our ROML formulation for multi-image object matching belongs to a more general class of MiAP for data association [13], with other vision applications such as multi-target tracking [19]. MiAP is proven to be NP-hard, and only implicit enumeration methods such as branch-and-bound are known to give an exact solution, which are however prohibitively slow for practical use. Classical approximate solution methods include greedy, Greedy Randomized Adaptive Search Procedure (GRASP) [51], and relaxation based methods [50].

Greedy approaches build a matching that has the lowest cost at each iteration, which has the obvious weakness that decisions once made, are fixed and may later be shown to be suboptimal. GRASP improves over greedy approaches by progressively constructing a list of best candidate matches and randomly selecting one from them. The process is repeated until all matches are exhausted. A final local search over the neighborhood of obtained matches may be used to further optimize the solution. In [50], Poore and Robertson presented a Lagrangian relaxation method that progressively relaxes the original and intermediate recovery MiAPs to linear assignment problems, by incorporating constraints of each MiAP into its objective function via the Lagrangian. However, this method involves implicitly enumerative procedure, and is difficult to implement and analyze.

Collins [19] recently proposed an iterated conditional modes (ICM) like method for video based multi-target tracking. His method is based on factoring the global decision variable for each target trajectory into a product of local variables defined for a target matching between each pair of adjacent frames. It then pair-wisely builds target matchings between adjacent frames by optimizing the corresponding local variables, but using a global cost function as matching criteria. However, the cost function in [19] is defined by enumerating for every possible target trajectory a constant-velocity motion energy, and the number of candidate trajectories grows exponentially with the number of frames. Both factors make it less applicable to the feature-based object matching problem considered in this paper. Nevertheless, our ROML formulation bears some spiritually similar idea with [19], in the sense that we also factor the global decision variables for multi-image feature matching into separate components. The difference is that we factor these global variables as a set of PPMs, each of which is to be optimized to identify inlier features from an image and re-arrange them in a proper order. We then treat the joint optimization of these PPMs as a regularized consensus problem in the context of distributed optimization [9], and solve it using a ADMM-based method [11]. By this means the original NP-hard combinatorial problem boils down as to iteratively solve a set of independent pair-wise matching problems, which turn out to be easily solved.

In the preparation of this paper, we notice that a related method called permutation synchronization [49] was recently proposed, which also addresses the MiAP by optimizing a permutation matrix for each feature set. However, [49] assumes initial matchings between each pair of the feature sets be available, and can only apply in the scenarios where there exist no outliers in each feature set, which make it less useful in the considered problem of feature-based object matching across a set of images.

3 Robust Object Matching Using Low-Rank and Sparse Constraints

Given a set of K images, we present in this section our problem formulation and algorithm for robust object matching. We consider the settings as stated in Problem 1. Assume n_k features $\{\mathbf{f}_i^k\}_{i=1}^{n_k}$ be extracted from the k^{th} image, where the feature vector $\mathbf{f}_i^k \in \mathbb{R}^d$ can be either image coordinates of the feature point, or region descriptors such as SIFT [42] that characterize the local appearance. It can also be a combination of them by

low-dimensional embedding [59]. In spite of these multiple choices, for now we generally refer to them as *features*. Discussion of different feature types and their applicable spectrums will be presented in Section 4. These n_k features are categorized as inliers or outliers. We assume at this moment that there are n inliers in each of the K images, where $n \leq n_k$ for $k = 1, \ldots, K$. We will discuss the case of missing inliers shortly. In such a setting every k^{th} image is represented as a set of n_k features, and the contained object instance is represented as the n inlier features.

3.1 Problem Formulation

Note that for inlier features in the K images, it is the feature similarity and geometric compatibility that determine they form an object *pattern* and this pattern repeats across the set of images. While similar outlier features may appear in multiple images, they just accidently do so in a random, unstructured way. Our formulation for object matching is essentially motivated by these observations. Concretely, denote $\overline{\mathbf{F}}^k = [\mathbf{f}_1^k, \ldots, \mathbf{f}_n^k] \in \mathbb{R}^{d\times n}$ as the matrix formed by inlier feature vectors in the k^{th} image, so defined are the matrices $\{\overline{\mathbf{F}}^1, \ldots, \overline{\mathbf{F}}^K\}$ for all the K images. Assume column vectors in each of these matrices are arrayed in the same order, i.e., inlier features in $\{\overline{\mathbf{F}}^1, \ldots, \overline{\mathbf{F}}^K\}$ are respectively corresponded, then the matrix formed by $\overline{\mathbf{D}} = [\mathrm{vec}(\overline{\mathbf{F}}^1), \ldots, \mathrm{vec}(\overline{\mathbf{F}}^K)] \in \mathbb{R}^{dn\times K}$ will be approximately low-rank, ideally rank one, where $\mathrm{vec}(\cdot)$ is an operator that vectorizes a matrix by concatenating its column vectors.

Now consider the general case that there are outliers. Denote $\mathbf{F}^k = [\mathbf{f}_1^k, \ldots, \mathbf{f}_{n_k}^k] \in \mathbb{R}^{d\times n_k}$ as the matrix having all n_k features of the k^{th} image as its columns, where feature vectors are placed in a random order. The matrices $\{\mathbf{F}^1, \ldots, \mathbf{F}^K\}$ for all K images are similarly defined. As aforementioned our interest for object matching is to identify the n inlier feature vectors from each matrix of $\{\mathbf{F}^1, \ldots, \mathbf{F}^K\}$, and establish correspondences among them. For any k^{th} image, this can be realized by the *partial permutation matrix* (PPM) defined by

$$\mathcal{P}^k = \{\mathbf{P}^k \in \mathbb{R}^{n_k\times n} \big| \mathbf{P}_{ij}^k \in \{0,1\}, \sum_i \mathbf{P}_{ij}^k = 1 \quad \forall j = 1, \ldots, n, \sum_j \mathbf{P}_{ij}^k \leq 1 \ \forall i = 1, \ldots, n_k\}. \quad (1)$$

Thus, there exist PPMs $\{\mathbf{P}^k \in \mathcal{P}^k\}_{k=1}^K$ such that inlier feature vectors are selected and corresponded in $\{\mathbf{F}^k\mathbf{P}^k \in \mathbb{R}^{d\times n}\}_{k=1}^K$, i.e., the matrix

$$\mathbf{D} = [\mathrm{vec}(\mathbf{F}^1\mathbf{P}^1), \ldots, \mathrm{vec}(\mathbf{F}^K\mathbf{P}^K)] \in \mathbb{R}^{dn\times K} \quad (2)$$

is rank deficient. Based on this low-rank assumption, feature correspondence can thus be formulated as the following problem to optimize $\{\mathbf{P}^k\}_{k=1}^K$

$$\min_{\{\mathbf{P}^k\in\mathcal{P}^k\}_{k=1}^K, \mathbf{L}} \mathrm{rank}(\mathbf{L}) \quad \text{s.t. } [\mathrm{vec}(\mathbf{F}^1\mathbf{P}^1), \ldots, \mathrm{vec}(\mathbf{F}^K\mathbf{P}^K)] = \mathbf{L}. \quad (3)$$

In practice, however, some inlier features characterizing the same local appearance information of object instances in different images could vary due to illumination change, object pose change, or other intra-category object variations. Some inlier features could also be missing due to partial occlusion of object instances. Thus the low-rank assumption used in (3) cannot be fully satisfied. To improve the robustness, we introduce a sparse error term into (3) to model all these contaminations of the data matrix $\mathbf{D}$, where we assume these errors only appear in a small fraction of $\mathbf{D}$. The formulation (3) can then be modified as

$$\min_{\{\mathbf{P}^k\in\mathcal{P}^k\}_{k=1}^K, \mathbf{L}, \mathbf{E}} \mathrm{rank}(\mathbf{L}) + \lambda\|\mathbf{E}\|_0 \quad \text{s.t. } [\mathrm{vec}(\mathbf{F}^1\mathbf{P}^1), \ldots, \mathrm{vec}(\mathbf{F}^K\mathbf{P}^K)] = \mathbf{L} + \mathbf{E}, \quad (4)$$

where $\|\cdot\|_0$ is ℓ_0-norm counting the number of nonzero entries, and $\lambda > 0$ is a parameter controlling the trade-off between rank of $\mathbf{L}$ and sparsity of $\mathbf{E}$.

3.2 The Algorithm

The optimization problem (4) is not directly tractable due to the following aspects: (1) both $\mathrm{rank}(\cdot)$ and $\|\cdot\|_0$ are non-convex, discrete-valued functions, minimization of which is NP-hard; (2) entries of $\{\mathbf{P}^k\}_{k=1}^K$ are constrained to be binary, resulting in a difficult nonlinear integer programming problem. To make it tractable, we first consider the recent convention of replacing $\mathrm{rank}(\cdot)$ and $\|\cdot\|_0$ with their convex surrogates $\|\cdot\|_*$ and $\|\cdot\|_1$ respectively [15], where $\|\cdot\|_*$ denotes nuclear norm (sum of the singular values) and $\|\cdot\|_1$ is ℓ_1-norm. Applying the same relaxation to (4) yields

$$\begin{aligned} \min_{\{\mathbf{P}^k\in\mathcal{P}^k\}_{k=1}^K, \mathbf{L}, \mathbf{E}} & \|\mathbf{L}\|_* + \lambda\|\mathbf{E}\|_1 \\ \text{s.t.} \quad & [\mathrm{vec}(\mathbf{F}^1\mathbf{P}^1), \ldots, \mathrm{vec}(\mathbf{F}^K\mathbf{P}^K)] = \mathbf{L} + \mathbf{E}, \\ & \mathcal{P}^k = \big\{\mathbf{P}^k \in \{0,1\}^{n_k\times n} \big| \mathbf{1}_{n_k}^\top\mathbf{P}^k = \mathbf{1}_n^\top, \\ & \qquad \mathbf{P}^k\mathbf{1}_n \leq \mathbf{1}_{n_k}\big\}, \ \forall\ k = 1, \ldots, K, \end{aligned} \quad (5)$$

where we have written the constraints of $\{\mathcal{P}^k\}_{k=1}^K$ in matrix form, and $\mathbf{1}_{n_k}$ (or $\mathbf{1}_n$) denotes a column vector of length n_k (or n) with all entry values of 1. We refer to the problem (5) as our framework of *Robust Object Matching using Low-rank and sparse constraints* (ROML).

The problem (5) involves jointly optimizing a set of K PPMs. As reviewed in Section 2, it is an instance of MiAP and proved to be NP-hard, for which approximate solution methods are practically used. To solve (5), note that it is a formulation of regularized consensus problem, where the local variables $\text{vec}(\mathbf{F}^k\mathbf{P}^k)$ (function of $\mathbf{P}^k$), $k = 1, \ldots, K$, are constrained to be equal to components (column vectors) of the global variable $\mathbf{L}+\mathbf{E}$, which is further regularized in the objective function. In literature, consensus problems are popularly solved using ADMM method in the context of distributed optimization [11,8,9]. The general ADMM method decomposes a global problem into local subproblems that can be readily solved. For consensus problems such as (5), ADMM decomposes optimization of $\mathbf{L}$, $\mathbf{E}$, and $\{\mathbf{P}^k\}_{k=1}^K$ into subproblems that update $\mathbf{L}$, $\mathbf{E}$, and each of $\{\mathbf{P}^k\}_{k=1}^K$ respectively. Thus joint optimization over $\{\mathbf{P}^k\}_{k=1}^K$ boils down as independent optimization of individual $\mathbf{P}^k$, $k = 1, \ldots, K$, in each ADMM iteration. However, the subproblem to update each $\mathbf{P}^k$ concerns with nonlinear integer programming. It is essential to understand the convergence property of ADMM under this condition, which we will discuss in Section 3.2.2 after presentation of our algorithmic procedure.

With $\mathbf{D} = [\text{vec}(\mathbf{F}^1\mathbf{P}^1), \ldots, \text{vec}(\mathbf{F}^K\mathbf{P}^K)]$, the augmented Lagrangian for (5) can be written as

$$\mathcal{L}_\rho(\mathbf{L}, \mathbf{E}, \{\mathbf{P}^k \in \mathcal{P}^k\}_{k=1}^K, \mathbf{Y}) = \|\mathbf{L}\|_* + \lambda\|\mathbf{E}\|_1 + \langle \mathbf{Y}, \mathbf{L}+\mathbf{E}-\mathbf{D}\rangle + \frac{\rho}{2}\|\mathbf{L}+\mathbf{E}-\mathbf{D}\|_F^2, \quad (6)$$

where $\mathbf{Y} \in \mathbb{R}^{dn\times K}$ is a matrix of Lagrange multipliers, ρ is a positive scalar, $\langle\cdot,\cdot\rangle$ denotes the matrix inner product, and $\|\cdot\|_F$ denotes the Frobenius norm. The ADMM algorithm iteratively estimates one of the matrices $\mathbf{L}$, $\mathbf{E}$, $\{\mathbf{P}^k\}_{k=1}^K$, and the Lagrange multiplier $\mathbf{Y}$ by minimizing (6), while keeping the others fixed. More specifically, our ADMM procedure consists of the following iterations

$$\mathbf{L}_{t+1} = \arg\min_{\mathbf{L}} \mathcal{L}_\rho\big(\mathbf{L}, \mathbf{E}_t, \{\mathbf{P}_t^k\}_{k=1}^K, \mathbf{Y}_t\big), \quad (7)$$

$$\mathbf{E}_{t+1} = \arg\min_{\mathbf{E}} \mathcal{L}_\rho\big(\mathbf{L}_{t+1}, \mathbf{E}, \{\mathbf{P}_t^k\}_{k=1}^K, \mathbf{Y}_t\big), \quad (8)$$

$$\{\mathbf{P}_{t+1}^k\}_{k=1}^K = \arg\min_{\{\mathbf{P}^k\in\mathcal{P}^k\}_{k=1}^K} \mathcal{L}_\rho\big(\mathbf{L}_{t+1}, \mathbf{E}_{t+1}, \{\mathbf{P}^k\}_{k=1}^K, \mathbf{Y}_t\big), \quad (9)$$

$$\mathbf{Y}_{t+1} = \mathbf{Y}_t + \rho\big(\mathbf{L}_{t+1} + \mathbf{E}_{t+1} - \mathbf{D}_{t+1}\big), \quad (10)$$

where t denotes the iteration number and we compute $\mathbf{D}_{t+1} = [\text{vec}(\mathbf{F}^1\mathbf{P}_{t+1}^1), \ldots, \text{vec}(\mathbf{F}^K\mathbf{P}_{t+1}^K)]$ after step (9).

The problems (7) and (8) for updating the global variables $\mathbf{L}$ and $\mathbf{E}$ are both convex programs. They can be explicitly written as the forms of the proximal operators associated with a nuclear norm or an ℓ_1-norm respectively [40]. To spell out the solutions, define the soft-thresholding operator for scalars as $\mathcal{T}_\tau[x] = \text{sign}(x)\cdot\max\{|x|-\tau, 0\}$, with $\tau > 0$. When applied to vectors or matrices, it operates element-wisely. The optimal solution to (7) and (8) can thus be written as

$$(\mathbf{U}, \mathbf{S}, \mathbf{V}) = \text{svd}\big(\mathbf{D}_t - \mathbf{E}_t - \frac{1}{\rho}\mathbf{Y}_t\big), \quad \mathbf{L}_{t+1} = \mathbf{U}\mathcal{T}_{\frac{1}{\rho}}\big[\mathbf{S}\big]\mathbf{V}^\top, \quad (11)$$

$$\mathbf{E}_{t+1} = \mathcal{T}_{\frac{\lambda}{\rho}}\big[\mathbf{D}_t - \mathbf{L}_{t+1} - \frac{1}{\rho}\mathbf{Y}_t\big]. \quad (12)$$

Optimization of (9) is more involved than those of (7) and (8), mostly because of the binary constraints enforced on the entries of $\{\mathbf{P}^k\}_{k=1}^K$. Plugging (6) into (9), we get an equivalent problem of (9) as

$$\min_{\{\mathbf{P}^k\in\mathcal{P}^k\}_{k=1}^K} \frac{\rho}{2}\big\|\mathbf{L}_{t+1} + \mathbf{E}_{t+1} + \frac{1}{\rho}\mathbf{Y}_t - [\text{vec}(\mathbf{F}^1\mathbf{P}^1), \ldots, \text{vec}(\mathbf{F}^K\mathbf{P}^K)]\big\|_F^2. \quad (13)$$

We observe that (13) can be decoupled into K independent subproblems, each of which concerns with optimization of one of the local variables $\{\mathbf{P}^k\}_{k=1}^K$. The k^{th} subproblem to update $\mathbf{P}^k$ is written as

$$\min_{\mathbf{P}^k\in\mathcal{P}^k} \frac{\rho}{2}\big\|\big(\mathbf{L}_{t+1} + \mathbf{E}_{t+1} + \frac{1}{\rho}\mathbf{Y}_t\big)\mathbf{e}_k - \text{vec}(\mathbf{F}^k\mathbf{P}^k)\big\|_2^2, \quad (14)$$

where $\mathbf{e}_k$ denotes a unit column vector with all entries set to 0 except the k^{th} one, which is set to 1. Denote $\theta^k = \text{vec}(\mathbf{P}^k) \in \mathbb{R}^{nn_k}$, $\mathbf{G}^k = \mathbf{I}_n \otimes \mathbf{F}^k \in \mathbb{R}^{dn\times nn_k}$, $\mathbf{J}^k = \mathbf{I}_n \otimes \mathbf{1}_{n_k}^\top \in \mathbb{R}^{n\times nn_k}$, $\mathbf{H}^k = \mathbf{1}_n^\top \otimes \mathbf{I}_{n_k} \in \mathbb{R}^{n_k\times nn_k}$, $\otimes$ is the Kronecker product, and $\mathbf{I}_n$ (or $\mathbf{I}_{n_k}$) is the identity matrix of size $n\times n$ (or $n_k\times n_k$). Using the fact $\text{vec}(\mathbf{XYZ}) = (\mathbf{Z}^\top\otimes\mathbf{X})\text{vec}(\mathbf{Y})$, we can rewrite (14) as the following equivalent problem to update θ^k

$$\min_{\theta^k} \frac{\rho}{2}\theta^{k\top}\mathbf{G}^{k\top}\mathbf{G}^k\theta^k - \mathbf{e}_k^\top\big[\mathbf{Y}_t^\top + \rho\big(\mathbf{L}_{t+1}+\mathbf{E}_{t+1}\big)^\top\big]\mathbf{G}^k\theta^k \quad \text{s.t. } \mathbf{J}^k\theta^k = \mathbf{1}_n,\ \mathbf{H}^k\theta^k \le \mathbf{1}_{n_k},\ \theta^k \in \{0,1\}^{nn_k}. \quad (15)$$

(15) appears to be a difficult integer constrained quadratic program. To solve it, a common approach is to relax the constraint set of (15) into its convex hull, and then project back the attained continuous-domain results by either thresholding or more complicated methods, which, however, cannot guarantee to get the optimal solution [38]. For the ROML problem (5), we assume that distinctive information of each column vector in any $\mathbf{F}^k$ of $\{\mathbf{F}^k\}_{k=1}^K$ is represented by the relative values of its elements, rather than their absolute magnitude. In other words, multiplying each feature vector by a scaling factor does not change the pattern of each feature. Based on this assumption, we prove that (15) is equivalent to a linear sum assignment problem [13].

Theorem 1 *For the ROML problem (5), assuming distinctive information of each column vector in any* $\mathbf{F}^k$ *of* $\{\mathbf{F}^k\}_{k=1}^K$ *is represented by the relative values of its elements, (15) is always equivalent to the following formulation of linear sum assignment problem*

$$\min_{\theta^k} -\mathbf{e}_k^\top \big[\mathbf{Y}_t^\top + \rho\big(\mathbf{L}_{t+1} + \mathbf{E}_{t+1}\big)^\top\big]\mathbf{G}^k\theta^k$$
$$\text{s.t. } \mathbf{J}^k\theta^k = \mathbf{1}_n,\ \mathbf{H}^k\theta^k \le \mathbf{1}_{n_k},\ \theta^k \in \{0,1\}^{nn_k}. \quad (16)$$

Proof We prove the equivalence by showing that, under the considered assumption for the ROML problem (5), the objective function of (15) is equivalent to a linear function, as written in (16), which together with the constraints of (16), turns out to be a formulation of LSAP. Denote $\mathbf{p}_i^k \in \mathbb{R}^{n_k}$, $i = 1, \dots, n$, as columns of PPM $\mathbf{P}^k$. From the definitions of $\mathbf{G}^k$ and θ^k, it is straightforward to show that

$$\mathbf{G}^k\theta^k = \mathrm{vec}(\mathbf{F}^k\mathbf{P}^k) = \big[(\mathbf{F}^k\mathbf{p}_1^k)^\top, \dots, (\mathbf{F}^k\mathbf{p}_n^k)^\top\big]^\top \quad (17)$$

Since $\mathbf{F}^k = [\mathbf{f}_1^k, \dots, \mathbf{f}_{n_k}^k] \in \mathbb{R}^{d\times n_k}$, and from the constraints of $\mathbf{P}^k$ (explicitly stated in (5)), it is clear that each subvector $\mathbf{F}^k\mathbf{p}_i^k$, $i = 1, \dots, n$, of (17) selects one column feature vector from $\mathbf{F}^k$, with a unique index from the set $\{1, \dots, n_k\}$. From (17) we also have

$$\theta^{k\top}\mathbf{G}^{k\top}\mathbf{G}^k\theta^k = \|\mathbf{G}^k\theta^k\|_2^2 = \sum_{i=1}^{n}\|\mathbf{F}^k\mathbf{p}_i^k\|_2^2. \quad (18)$$

In case of $n_k = n$, i.e., there exist no outliers in the considered feature-based object matching, (18) is equal to a constant value no matter what feasible $\mathbf{P}^k$ or θ^k is used. In the more general case of $n_k > n$, since information of each of the feature vectors $\mathbf{f}_i^k$, $i = 1, \dots, n_k$, is preserved by relative values of its elements, we can always normalize them so that they have an equal Euclidean norm, i.e., $\|\mathbf{f}_1^k\|_2 = \cdots = \|\mathbf{f}_{n_k}^k\|_2 = c_k$. And (18) is again equal to a constant value no matter what feasible $\mathbf{P}^k$ or θ^k is used. We thus finish the proof. □

The LSAP (16) can be exactly and efficiently solved using a rectangular-matrix variant of the Hungarian algorithm [13]. After solving K (16)-like problems for $k = 1, \dots, K$, we get the updates of $\{\theta_{t+1}^k\}_{k=1}^K$ and compute $\mathbf{D}_{t+1} = \big[\mathbf{G}^1\theta_{t+1}^1, \dots, \mathbf{G}^K\theta_{t+1}^K\big]$. The Lagrange multiplier matrix $\mathbf{Y}_{t+1}$ is then updated using (10). Our ADMM procedure iteratively performs the steps (7), (8), (9), and (10), until a specified stopping condition is satisfied. Normally, the primal and dual residuals can be used as the stopping criteria [2]. To improve the convergence, a common practice is to use a monotonically increasing sequence of $\{\rho_t\}$. We also adopt this strategy. The pseudocode of our algorithm is summarized in Algorithm 1 [3].

Algorithm 1: Solving ROML by ADMM

input : Feature vectors $\mathbf{F}^k = [\mathbf{f}_1^k, \dots, \mathbf{f}_{n_k}^k] \in \mathbb{R}^{d\times n_k}$ (normalized as $\|\mathbf{f}_1^k\|_2 = \cdots = \|\mathbf{f}_{n_k}^k\|_2 = c_k$ when there exist outliers), $k = 1, \dots, K$, the number n of inliers, weight $\lambda > 0$, and initialization of $\{\mathbf{P}_0^k \in \mathcal{P}^k\}_{k=1}^K$, $\mathbf{L}_0 = 0$, $\mathbf{E}_0 = 0$, $\mathbf{Y}_0 = 0$, and $\rho_0 > 0$.

```
while not converged do
    (U, S, V) = svd(D_t − E_t − (1/ρ_t) Y_t).
    L_{t+1} = U T_{1/ρ_t}[S] V^⊤.
    E_{t+1} = T_{λ/ρ_t}[D_t − L_{t+1} − (1/ρ_t) Y_t].
    for each k do
        let θ_t^k = vec(P_t^k), G^k = I_n ⊗ F^k,
        J^k = I_n ⊗ 1_{n_k}^⊤, H^k = 1_n^⊤ ⊗ I_{n_k}, solve the LSAP
        problem (16) to get the update θ_{t+1}^k.
    end
    D_{t+1} = [G^1 θ_{t+1}^1, ..., G^K θ_{t+1}^K].
    Y_{t+1} = Y_t + ρ_t(L_{t+1} + E_{t+1} − D_{t+1}).
    ρ_{t+1} ← ρ_t.
    t ← t + 1.
end
```

output: solution $\{\mathbf{P}_t^k\}_{k=1}^K$, $\mathbf{L}_t$, and $\mathbf{E}_t$.

3.2.1 Discussion of Solving ROML Using ADMM

Solving the ROML formulation (5) establishes n sets of consistent feature correspondences across the given K images. In other words, it aims to find n "good" ones out of the total $\frac{1}{n!}(\frac{n_k!}{(n_k-n)!})^K$ feasible solutions, assuming $n_1 = \cdots = n_k = \cdots = n_K$. As reviewed in Section 2, ROML belongs to the more general class of MiAP. To see how ROML relates to MiAP, we write the standard MiAP formulation [13] for the considered

[2] For the ROML problem (5), the primal residual is $\mathcal{R}_{pri.}^{t+1} = \mathbf{L}^{t+1} + \mathbf{E}^{t+1} - \mathbf{D}^{t+1}$, and the dual residuals are $\mathcal{R}_{dual,\mathbf{L}}^{t+1} = \rho(\mathbf{E}^t + \mathbf{D}^t - \mathbf{E}^{t+1} - \mathbf{D}^{t+1})$ (w.r.t. the variable $\mathbf{L}$) and $\mathcal{R}_{dual,\mathbf{E}}^{t+1} = \rho(\mathbf{D}^t - \mathbf{D}^{t+1})$ (w.r.t. the variable $\mathbf{E}$).

[3] We note that the obtained solution $\{\mathbf{P}^k\}_{k=1}^K$ by solving the ROML problem (5) using Algorithm 1 belongs to a group of equivalent solutions, since there is no constraint on the order of columns in any of $\{\mathbf{P}^k\}_{k=1}^K$. It is always easy to transform $\{\mathbf{P}^k\}_{k=1}^K$ to some canonical form, e.g., by permuting columns of each of $\{\mathbf{P}^k\}_{k=1}^K$ according to sorted image coordinates of feature points in any of the K images. Without loss of generality we assume the solution $\{\mathbf{P}^k\}_{k=1}^K$ given by Algorithm 1 has been transformed to some canonical form for ease of evaluation.

feature-based object matching problem as

$$\begin{aligned}
&\min_{\{z_{i_1 i_2 \cdots i_K}\}} \sum_{i_1=1}^{n_1} \sum_{i_2=1}^{n_2} \cdots \sum_{i_K=1}^{n_K} a_{i_1 i_2 \cdots i_K} z_{i_1 i_2 \cdots i_K} \\
\text{s.t.} &\sum_{i_2=1}^{n_2} \sum_{i_3=1}^{n_3} \cdots \sum_{i_K=1}^{n_K} z_{i_1 i_2 \cdots i_K} \leq 1 \ , \ i_1 = 1, \ldots, n_1 \\
&\sum_{i_1=1}^{n_1} \sum_{i_3=1}^{n_3} \cdots \sum_{i_K=1}^{n_K} z_{i_1 i_2 \cdots i_K} \leq 1 \ , \ i_2 = 1, \ldots, n_2 \\
&\qquad \vdots \\
&\sum_{i_1=1}^{n_1} \sum_{i_2=1}^{n_2} \cdots \sum_{i_{K-1}=1}^{n_{K-1}} z_{i_1 i_2 \cdots i_K} \leq 1 \ , \ i_K = 1, \ldots, n_K \\
&\sum_{i_1=1}^{n_1} \sum_{i_2=1}^{n_2} \cdots \sum_{i_K=1}^{n_K} z_{i_1 i_2 \cdots i_K} = n \ , \ z_{i_1 i_2 \cdots i_K} \in \{0, 1\}, \quad (19)
\end{aligned}$$

where $i_k \in \{1, \ldots, n_k\}$ indexes the n_k feature vectors extracted from the k^{th} image, $k = 1, \ldots, K$. The global decision variable $z_{i_1 i_2 \cdots i_K}$ is equal to 1 when the corresponding feature points are matched across the K images, with each feature from one of the K images, and $a_{i_1 i_2 \cdots i_K}$ denotes the cost of this matching. By factoring/reformulating the set of global decision variables $\{z_{i_1 i_2 \cdots i_K}\}$ as PPMs $\{\mathbf{P}^k\}_{k=1}^K$ defined by (1), we get the following equivalent problems

$$\min_{\{\mathbf{P}^k \in \mathcal{P}^k\}_{k=1}^K} \sum_{i=1}^{n} \sum_{i_1=1}^{n_1} \cdots \sum_{i_K=1}^{n_K} a_{i_1 \cdots i_K} \prod_{k=1}^{K} p_{i_k, i}^k, \quad (20)$$

$$\min_{\{\mathbf{P}^k \in \mathcal{P}^k\}_{k=1}^K} \sum_{i=1}^{n} \sum_{i_1=1}^{n_1} \cdots \sum_{i_K=1}^{n_K} \left\| [\mathbf{f}_{i_1}^1, \ldots, \mathbf{f}_{i_K}^K] \right\|_* \prod_{k=1}^{K} p_{i_k, i}^k, \quad (21)$$

$$\min_{\{\mathbf{P}^k \in \mathcal{P}^k\}_{k=1}^K} \sum_{i=1}^{n} \left\| [\mathbf{F}^1 \mathbf{p}_i^1, \ldots, \mathbf{F}^K \mathbf{p}_i^K] \right\|_*, \quad (22)$$

where $i \in \{1, \ldots, n\}$ indexes the n inlier matches, and $p_{i_k, i}^k$ and $\mathbf{p}_i^k$ denote the (i_k, i) entry and i^{th} column of $\mathbf{P}_k$ respectively. In (21) and (22), we have used nuclear norm of the matrix formed by a candidate match of K feature vectors as the cost coefficient $a_{i_1 \cdots i_K}$. As an instance of MiAP, jointly optimizing the set of PPMs in the above equivalent problems is NP-hard. Approximate methods are thus important to get practically meaningful solutions. In fact, due to inevitable noise in cost coefficients, e.g., that generated by various variations of object instances in different images, it is often sufficient to find suboptimal solutions that are within the noise level of the optimal one.

To understand how we have developed an approximate method in preceding sections, we slightly modify (22) by vertically arraying the n matrices $[\mathbf{F}^1 \mathbf{p}_i^1, \ldots, \mathbf{F}^K \mathbf{p}_i^K]$, $i = 1, \ldots, n$, as a bigger matrix, resulting in the following problem to optimize $\{\mathbf{P}^k\}_{k=1}^K$

$$\min_{\{\mathbf{P}^k \in \mathcal{P}^k\}_{k=1}^K} \left\| [\mathrm{vec}(\mathbf{F}^1 \mathbf{P}^1), \ldots, \mathrm{vec}(\mathbf{F}^K \mathbf{P}^K)] \right\|_*, \quad (23)$$

which turns out to be equivalent to a nuclear norm relaxed version of (3). Indeed, by introducing the global variable $\mathbf{L}$ in (3), and also the global variables $\mathbf{L}$ and $\mathbf{E}$ in (4) and (5) for a robust extension, we essentially formulate the multi-image object matching version of MiAP as a regularized consensus problem [9], with $\|\mathbf{L}\|_* + \lambda \|\mathbf{E}\|_1$ as the regularization term. It becomes well suited to be solved using distributed optimization methods such as ADMM. As in the presented ADMM procedure (7)-(10), the "fusion" steps (7) and (8) collect information of the t^{th} iteration $\{\mathbf{P}_t^k\}_{k=1}^K$ to update $\mathbf{L}$ and $\mathbf{E}$, and the "broadcast" step (9) independently updates each $\mathbf{P}^k$ of $\{\mathbf{P}^k\}_{k=1}^K$, using the updated fusion centers $\mathbf{L}_{t+1}$ and $\mathbf{E}_{t+1}$. Our proposed ADMM method thus belongs to a strategy of "fusion-and-broadcast", for ROML and the more general MiAP.

3.2.2 Convergence Analysis

The ADMM method is proven to converge to global optimum under some mild conditions for linearly constrained convex problem whose objective function is separable into two individual convex functions with non-overlapping variables (see [31] and references therein). In our case, the ROML problem (5) is nonconvex, due to the binary constraint associated with $\{\mathbf{P}^k\}_{k=1}^K$. The convergence property of ADMM for nonconvex problems is still an open question in theory. However, it is not uncommon to see that ADMM has served a powerful heuristic for some nonconvex problems in practice [64,56]. In the following, we present simulated experiments that demonstrate the excellent convergence property of ADMM for ROML.

Specifically, we generated synthetically $K = 30$ groups of vectors simulating extracted feature vectors from K images, with dimension of each vector $\mathbf{f}$ as $d = 50$. There were 50 vectors including both inliers and outliers in each group. The inliers were produced by randomly generating d-dimensional vectors whose entries were drawn from i.i.d. normal distribution, and were shared in each of the K groups. The outliers were similarly produced by randomly generating d-dimensional vectors following i.i.d. normal distribution, but were independently generated for each group. We then added sparse errors of large magnitude to both inlier and outlier vectors. For each vector $\mathbf{f}$, the error values were uniformly drawn from the range $[-2\max(\mathrm{abs}(\mathbf{f})), 2\max(\mathrm{abs}(\mathbf{f}))]$. Finally, we normalized all vectors to constant ℓ_2-norm to fit with our algorithmic settings.

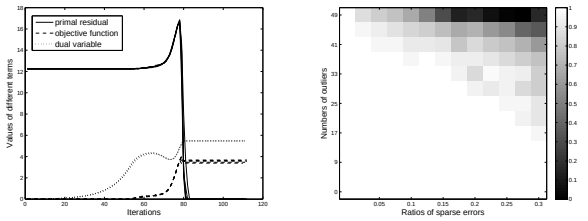

(a) (b)

Fig. 1 Simulation of Algorithm 1: (a) convergence plot in terms of the primal residual, objective function, and dual variable; (b) recovery precisions under varying numbers of outliers and ratios of sparse errors.

We investigated the convergence and recovery properties of our algorithm under varying numbers of outliers and ratios of sparse errors. The numbers of outliers in each group were ranged in $[0, 49]$, and the ratios of sparse errors in each vector were ranged in $[0, 0.3]$. We set $\lambda = 5/\sqrt{d}$ in these experiments. Denote the ground truth PPMs of any test setting as $\{\mathbf{P}^{k*}\}_{k=1}^{K}$, and the recovered PPMs as $\{\mathbf{P}^{k}\}_{k=1}^{K}$. The recovery precision is computed as $\sum_{k=1}^{K} \|\mathbf{P}^{k} \circ \mathbf{P}^{k*}\|_0 / \sum_{k=1}^{K} \|\mathbf{P}^{k*}\|_0$, where $\circ$ is Hadamard product. For each setting of outlier number and sparse error ratio, we run 5 random tests and averaged the results. Figure 1-(b) reports the recovery precisions under different settings, which shows that our algorithm works perfectly in a large range of outlier numbers and ratios of sparse errors. For one of them (the outlier number is 45 and sparse error ratio is 0.2), we plot in Figure 1-(a) its convergence of 5 random tests in terms of the primal residual ($\|\mathbf{L} + \mathbf{E} - \mathbf{D}\|_F$), objective function ($\|\mathbf{L}\|_* + \lambda\|\mathbf{E}\|_1$), and dual variable ($\|\mathbf{Y}\|_F$). Convergence properties under other settings are similar to Figure 1-(a).

3.2.3 Computational Complexity

For ease of analysis we assume here $n_1 = \cdots = n_k = \cdots = n_K > n$. Using an efficient implementation of the Hungarian algorithm [13], the complexity for solving the LSAP is $\mathcal{O}(n_k^3)$. The overall complexity for each iteration of Algorithm 1 is $\mathcal{O}(Kn_k^3 + Kdn^2n_k + K^2dn)$. The number of iterations for Algorithm 1 to converge depends on the initial value of ρ_0 and the factor at which ρ_t increases after each iteration. If ρ_t increases too fast, it has the risk of converging to worse local optima [40]. In our experiments, without mentioning we always set $\rho_0 = 1\mathrm{e}^{-6}$ and increase it iteratively with a factor of 1.01. Under this setting, it normally takes $500 \sim 1000$ iterations for Algorithm 1 to converge. In Section 5, we also report practical computation time of our method and compare with that of competing ones.

4 Choices of Feature Types and Their Applicable Spectrums

In the previous sections, we have represented an image as a set of features, where features generally refer to vectors characterizing image points and local regions centered on them. The task of object matching is then posed as Problem 1. Depending on different applications, these features can be chosen as either image coordinates, local region descriptors, or combination of them encoding both spatially structural and local appearance information. In the following, we present details of different choices of feature types and their applicable spectrums for robust object matching.

4.1 Image Coordinates

Given a set of points in an image, their coordinates can be directly used as features. In fact, coordinates of a set of inlier points in an image encode geometric relations among them, and it is the geometric structure of these points that determines the object pattern, and also provides a constraint for use in object matching. Image coordinates based features have been intensively used in early shape matching works [53, 55, 10, 17, 5].

For a moving rigid object in a video sequence or images of a rigid object captured from different viewpoints, denote $\mathbf{f}_i^k = [x_i^k, y_i^k]^\top \in \mathbb{R}^2$, $i = 1, \ldots, n_k$, as image coordinates based n_k features extracted from the k^{th} image. Let $\mathbf{F}^k = [\mathbf{f}_1^k, \ldots, \mathbf{f}_{n_k}^k] \in \mathbb{R}^{2\times n_k}$. It has been shown in [58] that the matrix, defined by

$$\mathbf{D}' = \left[(\mathbf{F}^1\mathbf{P}^1)^\top, \ldots, (\mathbf{F}^K\mathbf{P}^K)^\top\right]^\top \in \mathbb{R}^{2K\times n}, \tag{24}$$

is highly rank deficient (at most rank 4 when considering translation and there is no measurement noise), if correct PPMs $\{\mathbf{P}^k\}_{k=1}^K$ are used so that n inlier points can be selected from each of $\{\mathbf{F}^k\}_{k=1}^K$ and corresponding points $\{\mathbf{f}_i^k\}_{k=1}^K$, $i \in \{1, \ldots, n\}$, can be aligned in the same column of $\mathbf{D}'$. (24) is slightly different from the formation of $\mathbf{D}$ in (2). By applying the same low-rank and sparse constraints as in (5), we will show in Section 5.1 that image coordinates based features are very useful for matching rigid objects.

4.2 Local Region Descriptors

It is also straightforward to use region descriptors characterizing locally visual appearance information as features. These include SIFT [42], HOG [22], Geometric Blur [7, 6], GIST [46], or even raw pixels of local patches. In general, these feature descriptors have the properties of invariance and distinctiveness. The invariance

property makes it possible to match salient features extracted from images under geometric transformation or illumination change, while feature distinctiveness is important to differentiate between different salient regions. Features of such kind can be used in scenarios where they are discriminative enough for matching, or geometric constraints between feature points are not available, such as common object localization [23,66]. In Section 6, we present how ROML can be applied to this application.

4.3 Combination of Image Coordinates and Region Descriptors

Local region descriptors alone could be ambiguous for feature matching when there exist repetitive textures or less discriminative local appearance in images. To improve the matching accuracy, it is necessary to exploit the geometric structure of inlier points that consistently appears in each of the set of images. In literature, there are many ways to exploit such geometric constraints, such as pair-wise compatibility of feature correspondences used in graph matching [37,60], or linear-form constraints benefiting from a template image [39,32]. In this work, we consider a simple method introduced in [59]. For any interest point in each of the set of images, this method learns a low-dimensional embedded feature vector that combines information of both the local appearance and the spatial relations of this point relative to other points in the image. We present details of how to compute this type of learned embedded features in Appendix A.

As suggested by Theorem 1, when there are no outliers, the thus learned features can be directly used in our ROML framework. When there exist outliers in any of the set of images, we can always normalize those features to let them have constant ℓ_2-norm, and our method still applies. Since this type of learned features encode both appearance and spatial layout information, our method can potentially apply in more general settings, such as matching of non-rigid, articulated objects, or instances of a same object category. Experiments in Section 5 show the promise.

5 Experiments

In this section, we present experiments to show the effectiveness of ROML for robustly matching objects in a set of images. We consider different testing scenarios from the relatively simple rigid object matching, to the more challenging matching of object instances of a common category, and tracking a non-rigid object in a video sequence. For these testing scenarios, we choose appropriate feature types of either image coordinates or combination of image coordinates and local region descriptors, while features of region descriptors alone will be used in Section 6 for the application of common object localization. In the following experiments, without mentioning we always set the penalty parameter $\lambda = 5/\sqrt{dn}$ when solving the ROML problem (5) using Algorithm 1, where ρ was initially set as $1\mathrm{e}^{-6}$ and iteratively increased with a factor of 1.01.

5.1 Rigid Object with 3D Motion

The CMU "Hotel" sequence consists of 101 frames of a toy hotel building undergoing 3D motion [1]. Each frame has been manually labelled with the same set of 30 landmark points [14]. We use the "Hotel" sequence to show that ROML can be applied using image coordinates as features for matching rigid objects. In particular, we sampled 15 frames out of the total 101 frames (every 7 frames), in order to simulate the wide baseline matching scenario. Image coordinates of landmark points in these 15 frames were arranged into a matrix $\mathbf{D}' \in \mathbb{R}^{30\times 30}$ as defined in (24). We used Algorithm 1 to optimize a PPM for each frame, where the penalty parameter was set as $\lambda = 5/\sqrt{30}$, and ρ was initialized as $1\mathrm{e}^{-6}$ and iteratively increased with a factor of 1.0001.

We compare our method with representative pair-wise graph matching methods including Dual Decomposition (DD) [60], SMAC [21], and Learning Graph Matching (LGM) [14], which are based on either linear or quadratic assignment formulations, and also with more related methods [47,59] that are able to simultaneously match the set of 15 frames. For the former set of methods, matchings between a total of 105 frame pairs need to be established. Note that although all these methods are based on image coordinates, many of them have used the advanced shape context features [5]. To evaluate the performance of different methods, we use the match ratio criteria [4]. Table 1 reports the match ratios of different methods, where results of SMAC and LGM are from [60,59]. Table 1 tells that ROML and One-Shot [59] achieve the best performance (no matching error). However, One-Shot [59] uses shape context feature to characterize each landmark point, and it performs a low-dimensional embedding combining information of both geometric structure and local

[4] Denote n as the number of ground truth feature correspondences between an image pair, and $\bar{n}$ as the number of identified ground truth correspondences for this image pair. The match ratio is computed as $\frac{\sum \bar{n}}{\sum n}$, where $\sum$ sums over all image pairs.

Table 1 Results of different methods on the "Hotel" sequence. Accuracies are measured by the match ratio criteria.

Methods	DD [60]	SMAC [21]	LGM [14]	RankCon. [47]	One-Shot [59]	ROML-Pair	ROML
Accuracies	99.8%	84%	90%	57%	**100**%	**100**%	**100**%

Table 2 Match ratios of different methods on the "Hotel" sequence with varying numbers of missing inlier points in each frame.

No. of missing points	RankCon. [47]	One-Shot [59]	ROML
1	43%	76%	**95**%
3	26%	64%	**79**%
5	23%	59%	**71**%

descriptors of landmark points, while ROML just directly uses image coordinates. The method [47] also exploits low-rank constraints, however, its performance is much worse than that of ROML. It is probably due to the nuclear norm relaxation and ADMM optimization used in our framework. It is also interesting to observe that under the setting of pair-wise matching as for methods [60,21,14], ROML still performs perfectly on the "Hotel" sequence ("ROML-Pair" in Table 1). In fact, since the problem size of matching each frame pair is smaller, ROML converges faster with less iterations.

Compared to [47,59], ROML has the additional advantage of being more robust against missing inliers. To verify, we performed another experiment by removing randomly selected landmark points in each frame. For each removed landmark point, we also generated arbitrary image coordinates for it and made sure the generated coordinates were far enough away from the true ones, in order to fit with the algorithmic settings of these comparative methods. We set $\lambda = 2/\sqrt{30}$ in Algorithm 1. The One-Shot method [59] uses k-means clustering to obtain feature correspondences in the learned feature space. We chose its best-performing dimensionality of learned features, and run 10 trials of k-means clustering and averaged the results. Parameters of Rank Constraints [47] has also been tuned to its best performance. Table 2 reports the match ratio results, where matching accuracies are computed over non-missing points only. Table 2 clearly shows that ROML is less influenced when there exist missing inlier points.

5.2 Object Instances of a Common Category

In this section, we test how ROML performs to match object instances belonging to a same object category. We used 6 image sets of different categories from Caltech101 [28], MSRC [2], and the Internet. Numbers of images in these 6 sets ranged from 16 to 25. For each image, interest points were detected by SIFT: the numbers of detected interest points per image were from 27 to 174, out of which we manually labelled inlier points as matching ground truth. When some inlier points were not detected by SIFT in some images, we also manually labelled them in order to produce consistent sets of inlier points across the sets of images. In these experiments, we chose the low-dimensional embedded feature representation [59], as explained in Section 4.3, which encodes information of both geometric structure and descriptor similarity. To learn the embedded features, we used Geometric Blur descriptors [6] to characterize local regions around interest points, and Euclidean distances between points in each image for measuring geometric relations. The parameters for embedded feature learning were set as $\sigma_{spa.} = 10$ and $\sigma_{des.} = 0.2$ (cf. Appendix A for the definition of $\sigma_{spa.}$ and $\sigma_{des.}$), and the dimensionality of learned features was set as $d = 60$. Feature matching was realized by solving (5) using Algorithm 1.

We compare our method with One-Shot [59], and also with several recent graph matching methods including DD [60], RRWM [16], and SM [37], and hypergraph matching methods including TM [24], RRWHM [35], and ProbHM [62]. For One-Shot, we chose its best-performing dimensionality of the learned embedded features as our method used, and run 10 trials of k-means clustering and averaged the results. For pair-wise graph matching methods [60,16,37,24,35,62], we generated a total of $\frac{K!}{(K-2)!2!}$ image pairs for each test set with K images. These graph matching methods characterize interest points in each image by both their spatial relations and their respective local region descriptors. For a fair comparison, we used the same Geometric Blur region descriptors as our method used [5], when producing results using codes of these methods. Their parameters were also tuned to their respective best performance on the 6 test sets.

Table 3 reports results of different methods in terms of match ratio. Example feature correspondences for DD [60] and our method are shown in Figure 2. Table 3 and Figure 2 suggest that for the relatively simple

[5] We have also tried the SIFT descriptor [42] to characterize appearance of local regions around interest points. The matching accuracies using SIFT were slightly worse than those using Geometric Blur for both our method and these comparative methods.

"Airplane", "Motorbike", and "Face" image sets, our method gives very good matching results. The "Car", "Bus", and "Bank of America (BoA)" sets are more difficult due to the cluttered background, large viewpoint changes, or intra-category variations between different instances. Our method still gives reasonably good and consistent matching results. Both One-Shot and our method can match multiple images simultaneously. Our results are much better than those of One-Shot, which shows that One-Shot cannot perform well in the presence of outliers, and also that our ROML formulation optimized by the ADMM method is very effective for multi-image feature matching. Our method greatly outperforms graph and hyper-graph matching methods. It demonstrates that leveraging more object pattern constraints (i.e., geometric and feature similarity constraints) from multiple images is very useful for feature matching. Moreover, Figure 2 suggests that our matching results across the 4 images are more consistent than those from graph matching methods: another desired property for many computer vision applications. In Table 3, we also compare with our previous method [63] (Prev [63] in Table 3). Results of Prev [63] are obtained using the same low-dimensional embedded features as ROML does. Table 3 tells that on all the 6 image sets, matching accuracies of ROML are much better than those of Prev [63]. The improved accuracy comes from the new way of PPM optimization, i.e., exactly solving an equivalent LSAP in the present paper instead of sub-optimally solving two costly subproblems in [63], as we have discussed in Section 1.

Note that results of graph/hyper-graph matching methods reported in Table 3 are obtained by searching $\min(n_p, n_q)$ correspondences between any two images with n_p and n_q interest points respectively, and counting correct ones out of them, while our method only searches for $n < \min(n_p, n_q)$ correspondences. This number n of inliers for each image set is assumed given in our method. In practice, however, the true number of inliers is unknown. It is interesting to investigate how ROML performs when providing values of n to Algorithm 1 that are different from the true one. We conducted such experiments using the same 6 image sets, and results of matching accuracies are plotted in Figure 3. Figure 3 shows that results of ROML are relatively stable when giving different values of n to Algorithm 1, which suggests that performance of ROML is less influenced by the given knowledge.of inlier number. It is more desirable to have a mechanism to automatically estimate the true number of inliers. We are interested in pursuing this in future research.

Different choices of dimensionality d in low dimensional feature learning may influence our method's performance. In Figure 4-(a), we plot our matching accuracies with different choices of d on these 6 test sets. It shows that better results can generally be obtained when $d = 50 \sim 100$. It is expected that our method performs well only when the size of image sets (the K value) is relatively large. In Figure 4-(b), we plot results of our method on the 6 test sets with different choices of K. It shows that when $K > 10$, our method can stably get good results, which confirms that simultaneously matching a set of images is very useful for robust object matching.

Except for matching accuracy, one may also be interested in comparing matching efficiency of different methods. We have analyzed the computational complexity of our proposed method in Section 3.2.3. In Table 4, we report practical computation time of different methods for those experiments reported in Table 3. These experiments were conducted on an Intel Xeon CPU running at 2.8GHz, using Matlab implementation of different methods. Table 4 suggests that ROML is much more efficient than the best-performing graph matching method DD [60], and is on par with other graph/hyper-graph matching methods. One-Shot [59] is very fast, however, its accuracy (reported in Table 3) is not satisfactory. As an improved method of our previous work [63], ROML is much more efficient than [63] as well. The improved efficiency is again due to the new way of PPM optimization in the present paper. In spite of this improved efficiency, most of ROML's computation is still on solving LSAPs for updating the set of K PPMs (steps $5 \sim 7$ in Algorithm 1), which concerns with K independent subproblems and are fully parallelizable. When implementing the PPM optimization steps in parallel (ROML-Parallel in Table 4), efficiency of ROML is further improved.

5.3 Non-Rigid Object Tracking

Lastly, we treat object tracking in a video sequence as a feature matching problem, and show how ROML can be adapted for non-rigid object tracking. Given a video sequence with interest points detected in each frame, we label inlier points from those detected in the first frame. The task of object tracking is to track these inlier points, which are supposed to be on the object of interest, across the video frames. We used a 25-frame "Tennis" sequence and a 50-frame "Marple" sequence [12] for this testing scenario. For the first sequence, we used KLT tracker [43] to detect 100 interest points in each frame, and labelled 11 inlier points from those detected in the first frame. For the second one, we detected 150 interest points and labelled 14 inliers. To adapt our method to this tracking scenario, we simply

Table 3 Match ratios of different methods on 6 image sets of different object categories.

Methods	RRWM [16]	SM [37]	TM [24]	RRWHM [35]	ProbHM [62]	DD [60]	OneShot [59]	Prev [63]	ROML
Airplanes	28%	54%	17%	54%	32%	70%	65%	87%	**95%**
Face	40%	57%	26%	54%	14%	64%	61%	53%	**87%**
Motorbike	50%	46%	23%	58%	28%	73%	68%	89%	**95%**
Car	26%	39%	12%	23%	12%	51%	50%	59%	**74%**
Bus	13%	25%	24%	43%	18%	52%	44%	64%	**75%**
BoA	7%	12%	6%	15%	7%	12%	16%	35%	**70%**

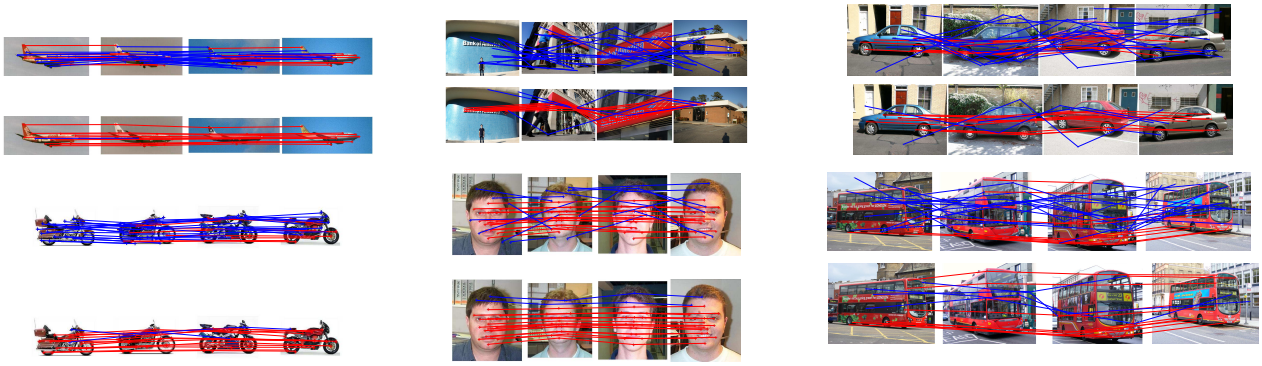

Fig. 2 Example feature correspondences among 4 images for different image sets. For every pair top is from DD [60], and bottom is from our method. Red lines represent identified ground truth correspondences, and blue lines are for false ones.

Table 4 Computation time (seconds) of different methods on the 6 image sets used in Table 3. All experiments were conducted on an Intel Xeon CPU running at 2.8GHz, using Matlab implementation of different methods.

Methods	Airplanes	Face	Motorbike	Car	Bus	BoA
RRWM [16]	22.92	211.10	53.20	260.83	157.30	240.60
SM [37]	5.44	61.71	17.61	39.31	20.61	43.15
TM [24]	44.01	418.30	145.38	282.92	187.39	315.04
RRWHM [35]	28.27	178.34	52.40	255.72	100.14	185.97
ProbHM [62]	39.58	375.48	128.45	252.04	168.70	269.31
DD [60]	2145.56	11714.50	3999.30	3688.99	3972.35	3827.95
One-Shot [59]	0.50	1.65	1.15	1.73	1.44	1.40
Prev [63]	227.87	1506.19	810.14	575.32	418.93	734.48
Prev [63]-Parallel	57.27	377.43	203.22	145.05	105.40	184.16
ROML	22.11	95.25	61.07	39.58	29.69	57.78
ROML-Parallel	6.79	26.79	17.99	13.78	9.40	17.32

fix $\mathbf{P}^1$ in steps 5 to 7 of Algorithm 1 so that it selects the n inliers labelled in the first frame, while optimizing the other $K-1$ PPMs $\{\mathbf{P}^k\}_{k=2}^K$. [6] We again used the type of learned embedded features as explained in Section 4.3. The parameter settings were the same as those used in Section 5.2.

We compare our method with a baseline KLT tracker, and recent graph and hyper-graph matching methods [60,16,37,24,35,62]. Since inlier points in the first frame are labelled ground truth, for graph and hyper-graph matching methods, we generated 24 and 49 frame pairs for the "Tennis" and "Marple" sequences respectively, i.e., between the first frame and each of the other frames, and used them for pair-wise matching. The other settings of these methods were the same as those used to produce Table 3 in Section 5.2. Parameters of these methods were also tuned to their respective best results on the two sequences. Table 5 reports the quantitative

[6] To make this scheme effective for tracking, we *emphasize* the labelled first frame by normalizing feature vectors of its interest points to have a larger value of ℓ_2-norm, compared to those in the other $K-1$ frames. This trick is important since when applying ROML to a video sequence with static background, it is possible that some of the interest points in the background are selected to form a "low-rank pattern", while those of the true pattern labelled in the first frame are treated as outliers. Theoretical correctness of this scheme is pending for proof. In practice, we observed that it worked well in the present tracking application. We defer the proof of this scheme in future research.

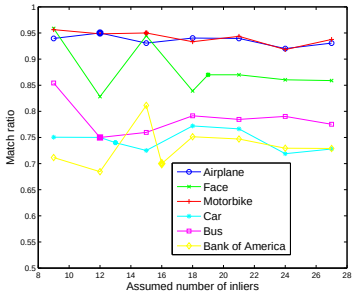

Fig. 3 Performance of ROML when providing different values of n (the assumed number of inliers for each image set) to Algorithm 1. Heavier marks correspond to the true numbers of inliers respectively for each image set.

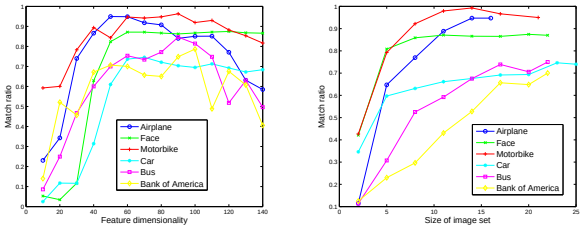

(a) (b)

Fig. 4 Performance of ROML using (a) different choices of dimensionality d in low-dimensional feature learning and (b) different sizes of image sets (the K values).

results of different methods in terms of match ratio. Example tracking/correspondences of interest points for DD [60] and our method are also shown in Figure 5. KLT tracker generally fails since there are abrupt motion and/or occlusion of inlier points in these two sequences. Compared to graph/hyper-graph matching methods, our method gives better results, which confirms the effectiveness of ROML for simultaneous multi-image object matching. We also compare with our previous method [63] in Table 5. Consistent to those results in Section 5.2, ROML again improves over Prev [63] on these two sequences.

Note that labelling inlier points in the first frame is essential in adapting our method to the object tracking application. When applying ROML to a video sequence with static background, without inlier labelling, it is possible that some of the interest points in the background are selected to form a "low-rank pattern", while the foreground object patterns are treated as outliers. We present in Figure 6 such a failure case of ROML using the same "Tennis" and "Marple" sequences. In fact, this difficulty of mining/matching foreground objects from video sequences with static background poses a common challenge for graph matching, hyper-graph matching, and our proposed methods.

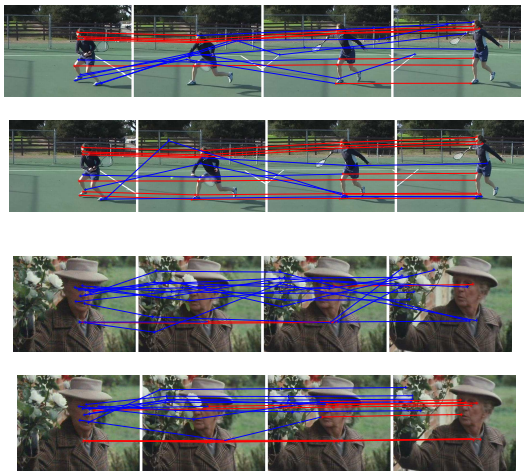

Fig. 5 Example tracking/correspondences of interest points among 4 frames of the "Tennis" and "Marple" sequences [12] respectively. For every pair top is from DD [60], and bottom is from our method. Red lines represent identified ground truth correspondences, and blue lines are for false ones.

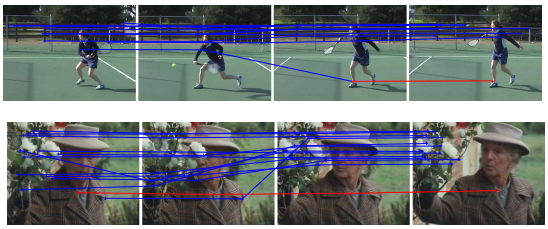

Fig. 6 Illustration of the failure of ROML on sequence data. Most of the identified correspondences by ROML are blue lines, which are, however, not the foreground objects of interest.

6 Common Object Localization

Learning models of object categories typically requires manually labelling a large amount of training images (e.g., up to a bounding box of the object of interest), which however, are expensive to obtain and may also suffer from unintended biases by annotators. A recently emerging research topic [61] considers automatically discovering and learning object models from a collection of unlabelled images. Given an image collection containing object instances belonging to unknown categories, the task is to identify the categories, localize object instances in images, and learn models for them so that the learned models can be applied to novel images for object detection. This is a weakly supervised (or unsupervised) learning scenario when the image collection is known to contain object instances of a single category (or multiple categories), which is in general ill-posed. A critical component for success of learning is precise

Table 5 Match ratios of different methods on the "Tennis" and "Marple" sequences [12].

Methods	KLT [43]	RRWM [16]	SM [37]	TM [24]	RRWHM [35]	ProbHM [62]	DD [60]	Prev [63]	ROML
Tennis	3%	23%	43%	13%	18%	16%	57%	52%	**73**%
Marple	4%	3%	25%	8%	13%	14%	23%	41%	**51**%

object localization inside each image. However, precise *common object localization* (COL) is extremely difficult given unknown object categories/models, and also large intra-category variations and cluttered background.

Many methods have been proposed for this challenging task in either weakly supervised or unsupervised settings [33,36,41,52]. Among them the methods [36,33,41] explicitly take object (or its associated parts/features) localization into account. These methods normally require the objects of interest covering a large portion of the images. More recently, saliency guided object learning techniques [66,23] are proposed, which exploit generic knowledge of "objectness" [3,25, 29] obtained from low-level image cues and/or learning from other irrelevant annotated images. Consequently, they can potentially locate object instances with large scale/appearance variations in cluttered background.

In this section, we present experiments to show how ROML can be applied to this COL task using local region descriptors as features. Similar to [23], we also sample candidate bounding boxes from each image based on their objectness scores, and use appropriate region descriptors to characterize the appearance inside each bounding box. We then optimize (5) to select a bounding box from each image, i.e., $n = 1$ for the PPMs to be optimized. Ideally the selected bounding boxes should localize object instances deemed common in the given image collection, i.e., the matrix $\mathbf{L}$ in (5) is rank deficient. We used the PASCAL datasets [27,26] for the COL experiments in both weakly supervised and unsupervised settings. For the weakly supervised case, we followed the same settings as in [23]. In particular, we used a subset of the PASCAL06 [27] train+val dataset containing all images of 6 classes (bicycle, car, cow, horse, motorbike, sheep) from the left and right viewpoints. We conducted COL on all images of each class-viewpoint combination, which are assumed to contain object instances of the same class at a similar viewpoint. To make the problem better defined, we followed [23] and removed images in which all objects are marked as difficult or truncated in the ground truth annotation. The PASCAL07 dataset [26] is more challenging as objects vary greatly in appearance, scale, and location. We also used 6 classes (aeroplane, bicycle, boat, bus, horse, and motorbike) of the PASCAL07 train+val dataset from the left and right viewpoints. The other settings were the same as for the PASCAL06 dataset. These classes of PASCAL06 and PASCAL07 datasets were chosen because they are the object classes on which fully supervised methods can perform reasonably well. For every image in one class-viewpoint combination, we used [3] to sample 100 bounding boxes proportionally to their probability of containing an object (the objectness score). To describe the region appearance inside each bounding box, we used the GIST descriptor with the default parameters as in [46], which gives a 512-dimensional feature vector. As suggested in [23], the shape and objectness score of a bounding box provide additional information that may help for COL. Let $\mathbf{f}$ be the GIST descriptor vector for a bounding box. To use shape and objectness score, we first augmented $\mathbf{f}$ with the aspect ratio ($width/height$) r of the bounding box, and then added perturbation noise $\mathbf{n} \in \mathbb{R}^{513}$ whose entries were drawn from normal distribution with standard deviation set as one minus objectness score of the bounding box. We used the thus produced vector $[\mathbf{f}^\top \ \kappa_r r]^\top + \kappa_{\mathbf{n}}\mathbf{n}$ as the feature for each sampled bounding box, where κ_r and $\kappa_{\mathbf{n}}$ are weighting parameters. We set $\kappa_r = 0.08$ and $\kappa_{\mathbf{n}} = 0.015$ for all the experiments reported in this section.

We measure COL performance by the percentage of correctly localized images out of all images in a class-viewpoint combination, where localization correctness in an image is based on PASCAL criteria, i.e., intersection of a bounding box with ground truth is more than half of their union. We compare with several baseline weakly supervised object localization and learning methods including MultiSeg [52] and Exemplar [18], and also with WSL-GK [23], which is saliency guided and performs EM-like alternation of localizing objects and learning the object class model. In the preparation of this paper, we notice that a more recent work [57] gives better COL performance by using more advanced saliency estimation method. Since this section is mainly to show the usefulness of ROML for the COL task, we will not pursue adopting this new saliency method to further improve our results.

Table 6 reports COL accuracies of different methods on the PASCAL06 and PASCAL07 datasets, which are obtained by averaging over all class-viewpoint combinations. Table 6 suggests that Objectness [3] gives very good initial candidates of object bounding boxes. Consequently, results of both our method and WSL-GK [23] on the PASCAL06 and PASCAL07 datasets compare

favorably with those from MultiSeg [52] and Exemplar [18]. For the PASCAL07 dataset, our method is comparable to WSL-GK [23] when no iterative steps of class learning are performed in [23], and greatly outperforms [23] for the PASCAL06 dataset, for which our result in fact approaches final result of [23], which is obtained after full steps of class learning and using richer feature representation including GIST, color information, and HOG for object shapes. Since the present paper is focusing on object matching and localization, we defer extension of our method for object class learning as future research.

We also conducted COL experiments in the unsupervised setting using 4 classes from the PASCAL06 (bicycle, car, cow, and sheep) and PASCAL07 (aeroplane, bus, horse, and motorbike) datasets respectively. Other data setups were the same as those in the above weakly supervised COL experiments. For either of the PASCAL06 and PASCAL07 datasets, we put all images of different classes from one viewpoint as an image collection, and applied ROML for object localization. Performance was again measured by the percentage of correctly localized images out of all images in a class-viewpoint combination. Table 7 reports detailed results of different class-viewpoint combinations, where we also list results of ROML in the weakly supervised setting. Table 7 tells that ROML performs consistently well in both weakly supervised and unsupervised object localization. Example images of these classes with localized bounding boxes are shown in Figure 7, where we also show the bounding boxes with the highest objectness score in each image and those of ROML in weakly supervised setting for comparison.

7 Conclusions

In this paper we propose a framework termed ROML, for robustly matching objects in a set of images. ROML is formulated as a rank and sparsity minimization problem to optimize a set of PPMs. The optimized PPMs identify inlier features from each image and establish their consistent correspondences across the image set. To solve ROML, we use the ADMM method, in which a subproblem associated with PPM updating is a difficult IQP. We prove that under widely applicable conditions, this IQP is equivalent to a formulation of LSAP, which can be efficiently solved by the Hungarian algorithm. Our framework is independent of the domain of applications and type of features. Experiments on feature based object matching, tracking, and common object localization show the efficacy of our proposed method. In future research, we are interested in applying/adapting ROML to other computer vision applications.

A Learning Features of Coordinates-Descriptor Combination

Local region descriptors alone could be ambiguous for feature matching when there exist repetitive textures or less discriminative local appearance in images. To improve the matching accuracy, it is necessary to exploit the geometric structure of inlier points that consistently appears in each of the set of images. In this work, we consider a simple method introduced in [59] to exploit such geometric constraints. The method derives an embedded feature representation that combines information of both spatial arrangement of feature points inside each image, and similarity of feature descriptors across images. We briefly summarize this method as follows.

Given a set of K images, denote $\mathbf{A}^k_{spa.} \in \mathbb{R}^{n_k \times n_k}$ as an affinity matrix that measures the spatial proximity of any two of the n_k extracted feature points in the k^{th} image, where spatial proximity can be either measured based on Euclidean distances of image coordinates of feature points, which is invariant to translation and rotation, or made affine invariant [59]. In this work, we compute $\mathbf{A}^k_{spa.}$ using Gaussian kernel as $\mathbf{A}^k_{spa.}(i,j) = \mathrm{e}^{-\|\mathbf{x}^k_i - \mathbf{x}^k_j\|^2/2\sigma^2_{spa.}}$, where $\mathbf{x}^k = [x^k, y^k]^\top$ denotes image coordinates in the k^{th} image, and $\sigma_{spa.}$ is a scaling parameter. Each feature point has an associated region descriptor. Denote $\mathbf{A}^{pq}_{des.} \in \mathbb{R}^{n_p \times n_q}$ as another affinity matrix, each entry of which measures the similarity of region descriptors between a pair of features selected from the p^{th} and q^{th} images respectively. $\mathbf{A}^{pq}_{des.}$ can be computed similar to $\mathbf{A}^k_{spa.}$ as $\mathbf{A}^{pq}_{des.}(i,j) = \mathrm{e}^{-\|\mathbf{f}^p_i - \mathbf{f}^q_j\|^2/2\sigma^2_{des.}}$, where $\mathbf{f}^p_i$ and $\mathbf{f}^q_j$ are feature descriptors from the p^{th} and q^{th} images respectively, and $\sigma_{des.}$ is a scaling parameter.

The method in [59] aims to learn embedded feature representations for all $N = \sum_{k=1}^K n_k$ points in the K images so that in the embedded space: (1) spatial structure of the point set in each image should be preserved; (2) features from different images with high descriptor similarity should be close to each other. Let $\{\mathbf{f}^k_i \in \mathbb{R}^d\}_{i=1}^{n_k}$, $k = 1, \ldots, K$, be the new features to be learned [7], the above objectives can be formalized as

$$\min \sum_{p,q} \sum_{i,j} \left\|\mathbf{f}^p_i - \mathbf{f}^q_j\right\|_2^2 \mathbf{A}^{pq}_{ij}, \tag{25}$$

where the matrix $\mathbf{A} \in \mathbb{R}^{N \times N}$ is defined as: $\mathbf{A}^{pq} = \mathbf{A}^k_{spa.}$ when $p = q = k$, $\mathbf{A}^{pq} = \mathbf{A}^{pq}_{des.}$ when $p \neq q$, and $\mathbf{A}^{pq} \in \mathbb{R}^{n_p \times n_q}$ is the (p,q) block of all the $K \times K$ blocks of $\mathbf{A}$. The objective function (25) turns to be a problem of Laplacian embedding [4]. Let $\tilde{\mathbf{F}} = [\mathbf{f}^1_1, \ldots, \mathbf{f}^1_{n_1}, \ldots, \mathbf{f}^K_1, \ldots, \mathbf{f}^K_{n_K}]^\top \in \mathbb{R}^{N \times d}$, (25) can be rewritten in matrix form as

$$\min_{\tilde{\mathbf{F}}} \operatorname{trace}\left(\tilde{\mathbf{F}}^\top \tilde{\mathbf{L}}_{\mathbf{A}} \tilde{\mathbf{F}}\right) \text{ s.t. } \tilde{\mathbf{F}}^\top \tilde{\mathbf{D}}_{\mathbf{A}} \tilde{\mathbf{F}} = \mathbf{I}, \tag{26}$$

where $\tilde{\mathbf{L}}_{\mathbf{A}} = \tilde{\mathbf{D}}_{\mathbf{A}} - \mathbf{A}$ is the Laplacian matrix of $\mathbf{A}$, and $\tilde{\mathbf{D}}_{\mathbf{A}}$ is a diagonal matrix with value of the i^{th} diagonal entry as $\sum_j \mathbf{A}_{ij}$. (26) is a generalized eigenvector problem: $\tilde{\mathbf{L}}_{\mathbf{A}}\mathbf{f} = \beta \tilde{\mathbf{D}}_{\mathbf{A}}\mathbf{f}$. Its optimal solution, i.e., the N new features in the d-dimensional embedded space, can be obtained by the bottom d nonzero eigenvectors.

[7] For consistency we use the same $\mathbf{f}$ for different feature types.

Table 6 COL accuracies of different methods on the PASCAL06 and PASCAL07 datasets. For objectness [3], sampled bounding box with the highest score in each image is considered as the estimated localization.

	Objectness [3]	MultiSeg [52]	Exemp. [18]	WSL-GK [23] (No Learning)	ROML	WSL-GK [23] (With Learning)
PASCAL06	51%	28%	45%	55%	**64**%	64%
PASCAL07	28%	22%	33%	**37**%	36%	50%

Table 7 COL accuracies of ROML for different class-viewpoint combinations of the PASCAL06 and PASCAL07 datasets in both weakly supervised and unsupervised settings.

	PASCAL06				PASCAL07			
	Bicycle	Car	Cow	Sheep	Aeroplane	Bus	Horse	Motorbike
Weakly Supervised - Left	84%	79%	60%	58%	26%	24%	40%	56%
Unsupervised - Left	80%	79%	60%	52%	30%	29%	35%	56%
Weakly Supervised - Right	69%	70%	66%	52%	38%	61%	35%	65%
Unsupervised - Right	67%	63%	57%	40%	28%	48%	41%	56%

Fig. 7 Example images with estimated bounding boxes of different class-viewpoint combinations from the PASCAL06 and PASCAL07 datasets. In each image, results of Objectness [3] with the highest objectness score (green box), and ROML in unsupervised (blue box) and weakly supervised (red box) settings are shown (they may coincide in some images where only one or two boxes are shown). Top part is for left viewpoint, and bottom part is for right viewpoint.

References

1. http://vasc.ri.cmu.edu//idb/html/motion/hotel/
2. The msrc dataset Http://research.microsoft.com/en-us / projects/objectclassrecognition/
3. Alexe, B., Deselaers, T., Ferrari, V.: What is an object? In: CVPR, pp. 73–80 (2010)
4. Belkin, M., Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation. Neural computation **15**(6), 1373–1396 (2003)
5. Belongie, S., Malik, J., Puzicha, J.: Shape matching and object recognition using shape contexts. IEEE Trans. on PAMI **24**(4), 509–522 (2002)
6. Berg, A.C., Berg, T.L., Malik, J.: Shape matching and object recognition using low distortion correspondences. In: CVPR, vol. 1, pp. 26–33 (2005)
7. Berg, A.C., Malik, J.: Geometric blur for template matching. In: CVPR, vol. 1, pp. I–607 (2001)
8. Bertsekas, D.P.: Nonlinear programming (1999)
9. Bertsekas, D.P., Tsitsiklis, J.N.: Parallel and distributed computation: numerical methods. Prentice-Hall, Inc. (1989)
10. Besl, P.J., McKay, N.D.: A method for registration of 3-d shapes. IEEE Trans. on PAMI **14**(2), 239–256 (1992)

11. Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J.: Distributed optimization and statistical learning via the alternating direction method of multipliers. Foundations and Trends® in Machine Learning **3**(1), 1–122 (2011)
12. Brox, T., Malik, J.: Object segmentation by long term analysis of point trajectories. In: ECCV, pp. 282–295 (2010)
13. Burkard, R.E., Dell'Amico, M., Martello, S.: Assignment Problems. Society for Industrial and Applied Mathematics (2009)
14. Caetano, T.S., McAuley, J.J., Cheng, L., Le, Q.V., Smola, A.J.: Learning graph matching. IEEE Trans. on PAMI **31**(6), 1048–1058 (2009)
15. Candès, E.J., Li, X., Ma, Y., Wright, J.: Robust principal component analysis? Journal of the ACM **58**(3), 11 (2011)
16. Cho, M., Lee, J., Lee, K.M.: Reweighted random walks for graph matching. In: ECCV, pp. 492–505 (2010)
17. Chui, H., Rangarajan, A.: A new point matching algorithm for non-rigid registration. Computer Vision and Image Understanding **89**(2), 114–141 (2003)
18. Chum, O., Zisserman, A.: An exemplar model for learning object classes. In: CVPR, pp. 1–8 (2007)
19. Collins, R.T.: Multitarget data association with higher-order motion models. In: CVPR, pp. 1744–1751 (2012)
20. Conte, D., Foggia, P., Sansone, C., Vento, M.: Thirty years of graph matching in pattern recognition. International Journal of Pattern Recognition and Artificial Intelligence **18**(03), 265–298 (2004)
21. Cour, T., Srinivasan, P., Shi, J.: Balanced graph matching. In: NIPS (2006)
22. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: CVPR, vol. 1, pp. 886–893 (2005)
23. Deselaers, T., Alexe, B., Ferrari, V.: Weakly supervised localization and learning with generic knowledge. IJCV **100**(3), 275–293 (2012)
24. Duchenne, O., Bach, F., Kweon, I.S., Ponce, J.: A tensor-based algorithm for high-order graph matching. IEEE Trans. on PAMI **33**(12), 2383–2395 (2011)
25. Endres, I., Hoiem, D.: Category independent object proposals. In: ECCV, pp. 575–588 (2010)
26. Everingham, M., Van Gool, L., Williams, C., Winn, L., Zisserman, A.: The pascal visual object classes challenge 2007
27. Everingham, M., Van Gool, L., Williams, C., Zisserman, A.: The pascal visual object classes challenge 2006
28. Fei-Fei, L., Fergus, R., Perona, P.: Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories. Computer Vision and Image Understanding **106**(1), 59–70 (2007)
29. Feng, J., Wei, Y., Tao, L., Zhang, C., Sun, J.: Salient object detection by composition. In: ICCV, pp. 1028–1035 (2011)
30. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM **24**(6), 381–395 (1981)
31. Goldfarb, D., Ma, S., Scheinberg, K.: Fast alternating linearization methods for minimizing the sum of two convex functions. Mathematical Programming **141**(1-2), 349–382 (2013)
32. Jiang, H., Yu, S.: Linear solution to scale and rotation invariant object matching. In: CVPR, pp. 2474–2481 (2009)
33. Kim, G., Faloutsos, C., Hebert, M.: Unsupervised modeling of object categories using link analysis techniques. In: CVPR, pp. 1–8 (2008)
34. Kuhn, H.W.: The hungarian method for the assignment problem. Naval research logistics quarterly **2**(1-2), 83–97 (1955)
35. Lee, J., Cho, M., Lee, K.M.: Hyper-graph matching via reweighted random walks. In: CVPR, pp. 1633–1640 (2011)
36. Lee, Y.J., Grauman, K.: Foreground focus: Unsupervised learning from partially matching images. IJCV **85**(2), 143–166 (2009)
37. Leordeanu, M., Hebert, M.: A spectral technique for correspondence problems using pairwise constraints. In: ICCV, vol. 2, pp. 1482–1489 (2005)
38. Leordeanu, M., Hebert, M., Sukthankar, R.: An integer projected fixed point method for graph matching and map inference. In: NIPS, vol. 2, p. 7 (2009)
39. Li, H., Huang, J., Zhang, S., Huang, X.: Optimal object matching via convexification and composition. In: ICCV, pp. 33–40 (2011)
40. Lin, Z., Chen, M., Ma, Y.: The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices. arXiv preprint arXiv:1009.5055 (2010)
41. Liu, D., Chen, T.: Unsupervised image categorization and object localization using topic models and correspondences between images. In: ICCV, pp. 1–7 (2007)
42. Lowe, D.G.: Object recognition from local scale-invariant features. In: ICCV, vol. 2, pp. 1150–1157 (1999)
43. Lucas, B.D., Kanade, T., et al.: An iterative image registration technique with an application to stereo vision. In: IJCAI, vol. 81, pp. 674–679 (1981)
44. Maciel, J., Costeira, J.P.: A global solution to sparse correspondence problems. IEEE Trans. on PAMI **25**(2), 187–199 (2003)
45. Mikolajczyk, K., Schmid, C.: Scale & affine invariant interest point detectors. IJCV **60**(1), 63–86 (2004)
46. Oliva, A., Torralba, A.: Modeling the shape of the scene: A holistic representation of the spatial envelope. IJCV **42**(3), 145–175 (2001)
47. Oliveira, R., Costeira, J., Xavier, J.: Optimal point correspondence through the use of rank constraints. In: CVPR, vol. 2, pp. 1016–1021 (2005)
48. Oliveira, R., Ferreira, R., Costeira, J.P.: Optimal multiframe correspondence with assignment tensors. In: ECCV, pp. 490–501 (2006)
49. Pachauri, D., Kondor, R., Singh, V.: Solving the multiway matching problem by permutation synchronization. In: NIPS, pp. 1860–1868 (2013)
50. Poore, A.B., Robertson III, A.J.: A new lagrangian relaxation based algorithm for a class of multidimensional assignment problems. Computational Optimization and Applications **8**(2), 129–150 (1997)
51. Robertson, A.J.: A set of greedy randomized adaptive local search procedure (grasp) implementations for the multidimensional assignment problem. Computational Optimization and Applications **19**(2), 145–164 (2001)
52. Russell, B.C., Freeman, W.T., Efros, A.A., Sivic, J., Zisserman, A.: Using multiple segmentations to discover objects and their extent in image collections. In: CVPR, vol. 2, pp. 1605–1614 (2006)
53. Scott, G.L., Longuet-Higgins, H.C.: An algorithm for associating the features of two images. Proc. of the Royal Society of London. Series B: Biological Sciences **244**(1309), 21–26 (1991)
54. Serre, T., Wolf, L., Poggio, T.: Object recognition with features inspired by visual cortex. In: CVPR, vol. 2, pp. 994–1000 (2005)
55. Shapiro, L.S., Michael Brady, J.: Feature-based correspondence: an eigenvector approach. Image and Vision Computing **10**(5), 283–288 (1992)

56. Shen, Y., Wen, Z., Zhang, Y.: Augmented lagrangian alternating direction method for matrix separation based on low-rank factorization. Optimization Methods and Software **29**(2), 239–263 (2014)
57. Siva, P., Russell, C., Xiang, T., Agapito, L.: Looking beyond the image: Unsupervised learning for object saliency and detection. In: CVPR, pp. 3238–3245 (2013)
58. Tomasi, C., Kanade, T.: Shape and motion from image streams under orthography: a factorization method. IJCV **9**(2), 137–154 (1992)
59. Torki, M., Elgammal, A.: One-shot multi-set non-rigid feature-spatial matching. In: CVPR, pp. 3058–3065 (2010)
60. Torresani, L., Kolmogorov, V., Rother, C.: Feature correspondence via graph matching: Models and global optimization. In: ECCV, pp. 596–609 (2008)
61. Tuytelaars, T., Lampert, C.H., Blaschko, M.B., Buntine, W.: Unsupervised object discovery: A comparison. IJCV **88**(2), 284–302 (2010)
62. Zass, R., Shashua, A.: Probabilistic graph and hypergraph matching. In: CVPR, pp. 1–8 (2008)
63. Zeng, Z., Chan, T.H., Jia, K., Xu, D.: Finding correspondence from multiple images via sparse and low-rank decomposition. In: ECCV, pp. 325–339 (2012)
64. Zhang, Y.: An alternating direction algorithm for nonnegative matrix factorization. Rice Univ. CAAM Tech. Rep. (2010)
65. Zhou, F., De la Torre, F.: Factorized graph matching. In: CVPR, pp. 127–134 (2012)
66. Zhu, J.Y., Wu, J., Wei, Y., Chang, E., Tu, Z.: Unsupervised object class discovery via saliency-guided multiple class learning. In: CVPR, pp. 3218–3225 (2012)